

# A note on Bayesian logistic regression for spatial exponential family Gibbs point processes

Tuomas A. Rajala

Department of Mathematical Sciences  
Chalmers University of Technology and University of Gothenburg  
S-41296 Gothenburg, Sweden  
`tuomas.rajala@chalmers.se`

November 4, 2014

## Abstract

Recently, a very attractive logistic regression inference method for exponential family Gibbs spatial point processes was introduced. We combined it with the technique of quadratic tangential variational approximation and derived a new Bayesian technique for analysing spatial point patterns. The technique is described in detail, and demonstrated on numerical examples.

*Keywords:* Exponential family model; Logistic regression; Bayesian inference

## 1 Introduction

Gibbs point processes are a popular class of models for interacting events taking place in spatial locations. Parametric inference for these inherently non-independent event models is costly due to analytically intractable normalizing constants in the likelihood. However, due to the recent development of approximate techniques such as the logistic regression technique described by [Baddeley et al. \(2014\)](#), the act of modeling is becoming free from the constraints of computational limitations.

In this paper we will describe a Bayesian version of the logistic regression technique of [Baddeley et al. \(2014\)](#). Bayesian methodology is not generally shunned upon by point pattern statisticians, so the slow growth in popularity amongst scientists is instead more likely a result of the need to employ MCMC algorithms (see overviews in e.g. [Gelfand et al., 2010](#); [Møller and Waagepetersen, 2007](#)). MCMC for point process models requires strong expertise in designing the algorithms even for the simplest models, and usually additional, costly simulation loops are needed to evaluate the chains' transitions probabilities. The proposed method requires no simulations and provides posterior distributions in a fraction of the running time of an MCMC algorithm.

This text is based on the setup described in detail by [Baddeley et al. \(2014\)](#), and we often refer to that paper to minimize repetition. Section 2 describes the details of the Bayesian version, and Section 3 depicts some examples on how the Bayesianity can be exploited for a flexible analysis of point pattern data.

## 2 Method

### 2.1 Model family and the approximative likelihood

We assume that we observe a stationary point process in some finite window  $W \subset \mathbb{R}^d$ . We denote the volume of the window by  $V := |W|$ , and write  $\omega \subset W$  for the observed, finite configuration of point locations  $\omega_i \in W$ . The cardinality is denoted by  $n = n(\omega) := |\omega \cap W|$ . We write  $u \in W$  for an arbitrary point in the window. The notation is slightly changed from that of [Baddeley et al. \(2014\)](#) to simplify the exposition of the model fitting algorithm later on.

The exponential family Gibbs models are defined through a density with respect to a unit rate Poisson process in  $\mathbb{R}^d$ . In the so called canonical parametrisation, a model of the exponential Gibbs model family has a parametric

density of the form

$$\tilde{f}(\omega|\theta) = \alpha(\theta)^{-1} H(\omega) \exp[\theta^T t(\omega)] \quad (1)$$

where  $\theta \in \mathbb{R}^p$  is a vector of parameters,  $t(\omega) = [t_1(\omega) \dots t_p(\omega)]^T \in \mathbb{R}^p$  is a vector of canonical statistics, function  $H(\omega)$  is a baseline or a reference density (mainly to have a hard-core effect), and  $\alpha(\theta)$  is the normalizing constant.

The normalizing constant  $\alpha(\theta)$  is usually not tractable, hindering direct likelihood based inference. Approximate likelihood inference for Gibbs models is usually based on local dynamics of the process described by the Papangelou conditional intensity, defined at any location  $u \in W$  as

$$\lambda(u; \omega) := \frac{\tilde{f}(\omega \cup u)}{\tilde{f}(\omega \setminus u)} = H(u; \omega) \exp[\theta^T t(u; \omega)] \quad (2)$$

with the convention  $\omega \setminus u = \omega$  if  $u \notin \omega$ . Here  $H(u; \omega) := 1\{H(\omega) > 0\} H(\omega \cup u) / H(\omega \setminus u)$  and  $t(u; \omega) := t(\omega \cup u) - t(\omega \setminus u)$  are the contributions of the point  $u$  to the density.

The logistic likelihood approximation described by [Baddeley et al. \(2014\)](#) replaces the true likelihood (1) with an unbiased estimation equation formally equal to that of standard logistic regression. The method is based on an auxiliary point configuration: First, we simulate a set of dummy locations  $d = \{u_j \in W : j = 1, \dots, m\}$  from a known process with a known intensity function  $\varrho(u) > 0$ . Then we concatenate the data and the dummy points to a configuration  $\bar{\omega} = \omega \cup d$ , keeping the order, and define a vector of data indicators  $y_i := 1\{u_i \in \omega\}$  for each  $u_i \in \bar{\omega}$ ,  $i = 1, \dots, n + m$ .

Then the likelihood (1) is well approximated by a logistic regression density of the form

$$f(\omega|\theta) = \prod_{i=1}^{n+m} p_i^{y_i} (1 - p_i)^{1-y_i} \quad (3)$$

where

$$p_i = \frac{\exp[\theta^T t(u_i; \omega) + o_i]}{1 + \exp[\theta^T t(u_i; \omega) + o_i]}$$

and  $o_i := \log[H(u_i; \omega) / \varrho(u_i)]$  are offset terms. Spatial trend components with possible location and covariate dependencies can be included as usual for logistic regression.

The power of the approach is that the original, computationally hard point process inference problem is approximated by much simpler standard exponential family form which can be solved in the generalized linear model (GLM) framework with standard software. In what follows a variational Bayes method for computing the solution is described.

## 2.2 Bayesian inference for logistic regression using variational approximation

The following variational technique for logistic regression was first described in 1996 and further developed by [Jaakkola and Jordan \(2000\)](#), motivated by requirements of fast learning of graphical models. It is based on a tight tangential bound for the function  $\log(1 + e^x)$ , and with Gaussian priors for  $\theta$  the technique leads to tractable Gaussian posteriors. The technique is part of a large group of variational techniques developed for Bayesian data analysis, often grouped under the name Variational Bayes (VB) techniques (see e.g. [Ormerod and Wand \(2010\)](#) for an overview).

Write  $N := n + m$ . To conform to the usual GLM notation, we will now write  $X = [\mathbf{x}_1 \dots \mathbf{x}_N]^T$  for the data matrix with the canonical statistics forming the rows,  $\mathbf{x}_i = t(u_i; \omega)$ . Using vector notation the logarithm of (3) becomes

$$\log f(\mathbf{y}|\theta) = \sum_{i=1}^{n+m} \{y_i(\mathbf{x}_i^T \theta + o_i) - \log[1 + \exp(\mathbf{x}_i^T \theta + o_i)]\} = \mathbf{y}^T (X\theta + \mathbf{o}) - 1_N^T \log[1 + \exp(X\theta + \mathbf{o})] \quad (4)$$

with  $1_N = [1 \dots 1]^T$  denoting the  $N$  length vector of 1's and  $\mathbf{o} = [o_1 \dots o_N]^T$ . The formulation in [Ormerod and Wand \(2010\)](#) has now been extended to include the offset terms.

The logarithm terms in (4) make it hard to derive posteriors for  $\theta$  as the normalizing integral is not tractable. The method of [Jaakkola and Jordan \(2000\)](#) works around this by approximating the logarithm terms with tangential lower bounds of quadratic form. The bound is based on the inequality

$$-\log[1 + \exp(\mathbf{x}_i^T \theta + o_i)] \geq \lambda(\xi_i)(\mathbf{x}_i^T \theta + o_i)^2 - \frac{1}{2}(\mathbf{x}_i^T \theta + o_i) + \gamma(\xi_i) \quad \forall \xi_i > 0$$

where the additional terms are  $\lambda(\xi_i) = -\tanh(\xi_i/2)/4\xi_i$  and  $\gamma(\xi_i) = \xi_i/2 - \log(1 + e^{\xi_i}) + \frac{\xi_i}{4} \tanh(\xi_i/4)$ . The  $\xi_i$ 's are called the variational parameters, and they determine the goodness of the approximation. Equality holds when  $\xi_i^2 = (\mathbf{x}_i^T \theta + o_i)^2$ . In vector form, applying the bound to each of the logarithm terms in (4) leads to the quadratic form

$$\begin{aligned} -1_N^T \log[1 + \exp(X\theta + \mathbf{o})] &\geq (X\theta + \mathbf{o})^T \Lambda(\xi)(X\theta + \mathbf{o}) - \frac{1}{2}1_N^T(X\theta + \mathbf{o}) + 1_N^T \gamma(\xi) \\ &= \theta^T X^T \Lambda(\xi) X \theta + 2\mathbf{o}^T \Lambda(\xi) X \theta + \mathbf{o}^T \Lambda(\xi) \mathbf{o} - \frac{1}{2}1_N^T(X\theta + \mathbf{o}) + 1_N^T \gamma(\xi) \end{aligned}$$

where  $\Lambda(\xi) = \text{diag}(\lambda(\xi_i))$ .

By setting a Gaussian prior  $\theta \sim N(\mu_0, \Sigma_0)$  the log of the joint density  $f(\mathbf{y}, \theta)$  in (4) has the closed form lower bound

$$\begin{aligned} \log \underline{f}(\mathbf{y}, \theta; \xi) &= -\frac{1}{2}\theta^T [\Sigma_0^{-1} - 2X^T \Lambda(\xi) X] \theta \\ &+ [(\mathbf{y} - \frac{1}{2}1_N + 2\Lambda(\xi)\mathbf{o})^T X + \mu_0^T \Sigma_0^{-1}] \theta \\ &+ 1_N^T \gamma(\xi) + \mathbf{o}^T \Lambda(\xi) \mathbf{o} + (\mathbf{y} - \frac{1}{2}1_N)^T \mathbf{o} - \frac{p}{2} \log(2\pi) - \frac{1}{2} \log |\Sigma_0| - \frac{1}{2} \mu_0^T \Sigma_0^{-1} \mu_0. \end{aligned}$$

The bound is proportional to a Gaussian density in  $\theta$  with posterior covariance matrix and mean given by

$$\begin{aligned} \Sigma_\xi^{-1} &= \Sigma_0^{-1} - 2X^T \Lambda(\xi) X \\ \mu_\xi &= \Sigma_\xi [(\mathbf{y} - \frac{1}{2}1_N + 2\Lambda(\xi)\mathbf{o})^T X + \mu_0^T \Sigma_0^{-1}]^T \end{aligned}$$

As these can be computed, we have derived a closed form for the (approximate) posterior distribution.

One detail remains: Both posterior parameters depend on the variational parameters which we need to define. A natural criterion for optimizing the variational parameters is given by the difference between the lower bound of the evidence  $\underline{f}(\mathbf{y}; \xi)$  and the true evidence  $f(\mathbf{y})$ , i.e. we aim to minimize the Kullback-Leibler divergence between the true evidence and the approximation. Due to Gaussianity of  $\theta$  the lower bound of the log-evidence can be solved in closed form as

$$\log \underline{f}(\mathbf{y}; \xi) = \frac{1}{2} \log |\Sigma_\xi| - \frac{1}{2} \log |\Sigma_0| + 1_N^T \gamma(\xi) + \frac{1}{2} \mu_\xi^T \Sigma_\xi^{-1} \mu_\xi - \frac{1}{2} \mu_0^T \Sigma_0^{-1} \mu_0 + \mathbf{o}^T \Lambda(\xi) \mathbf{o} + (\mathbf{y} - \frac{1}{2}1_N)^T \mathbf{o} \quad (5)$$

Numerical techniques could now be used for finding  $\xi$  that maximizes (5), but [Jaakkola and Jordan \(2000\)](#) proposed to use a simple expectation-maximization (EM) algorithm instead. With some matrix calculus, especially the linearity of traces, it can be shown that the E-step function has the form

$$Q(\xi' | \xi) = E_{\theta | \mathbf{y}; \xi} \log \underline{f}(\mathbf{y}, \theta; \xi') = \text{tr} \{ [X(\Sigma_\xi + \mu_\xi \mu_\xi^T) X^T + \mathbf{o} \mathbf{o}^T + 2X \mu_\xi \mathbf{o}^T] \Lambda(\xi') \} + 1_N^T \gamma(\xi') + \text{const.}$$

As  $\lambda(\xi_i)$  is a strictly increasing function of  $\xi_i > 0$ , the optimal values are given by the equation

$$\xi^2 = \text{diag}[X(\Sigma_\xi + \mu_\xi \mu_\xi^T) X^T + \mathbf{o} \mathbf{o}^T + 2X \mu_\xi \mathbf{o}^T]$$

cf. [Ormerod and Wand \(2010\)](#) with the current extension of including the offset terms.

A summary of the technique is given in Algorithm 1. The computational complexity of the technique differs from that of [Baddeley et al. \(2014\)](#) only in the logistic regression estimation part so we comment on the complexity of the variational approximation. The variational parameter vector is of size  $N = n + m$ , and each iteration the inverse and determinant operations are at most of order  $O(p^3)$ . Since usually  $p \ll N$  the complexity of the VB algorithm is dominated by the matrix multiplication with complexity of  $O(N^2 p)$ . In comparison, the complexity of solving iterated weighted least squares equation, the default technique in the `glm`-function used by `spatstat`-package, is of the same order.

---

**Algorithm 1** Posterior estimation for exponential family Gibbs point processes

---

```

1: Set  $\mu_0, \Sigma_0, H$  and  $\varrho$ 
2: Generate dummy points  $d$ 
3: Create vector  $\mathbf{y}$  and the matrix  $X$ 
4: Initialize  $\xi$ ,  $N \times 1$  non-negative vector
5: while log-evidence is increasing do
6:   Update  $\Sigma_\xi^{-1} = \Sigma_0^{-1} - 2X^T \Lambda(\xi) X$ 
7:   Update  $\mu_\xi = \Sigma_\xi[(\mathbf{y} - \frac{1}{2} \mathbf{1}_N + 2\Lambda(\xi)\mathbf{o})^T X + \mu_0^T \Sigma_0^{-1}]^T$ 
8:   Update  $\xi = \sqrt{\text{diag}[X(\Sigma_\xi + \mu_\xi \mu_\xi^T)X^T + \mathbf{o}\mathbf{o}^T + 2X\mu_\xi \mathbf{o}^T]}$ 
9: end while
10: Return  $(\mu_\xi, \Sigma_\xi)$ 

```

---

### 3 Examples

#### 3.1 Strauss process simulations

To assess the precision of the VB technique we compare it to the default GLM-fitter of the statistical software environment R. The comparison consists of fits to simulations from a range of homogeneous Strauss processes, a model family with repulsive interaction between events' locations. With some fixed parameter  $R$  denoting the interaction range, a Strauss process has the canonical statistics in the conditional intensity (2) of the form

$$t(u; \omega) = [1, n(\omega \cap b(u, R))]^T$$

where  $b(u, R)$  is the ball of radius  $R$  centered at  $u$ . The model has two parameters to estimate, the intensity related parameter  $\theta_1$  and the strength of interaction parameter  $\theta_2$ . Strict definition of the model requires  $\theta_2 < 0$ . Note that usually the model is parameterized by  $[\beta, \gamma]^T := \exp(\theta)$ . The range parameter  $R$  is not of canonical form, so the models are conditional on fixed  $R$ .

We simulated 100 realisations of the Strauss model with design  $\exp(\theta) \in \{100, 1000\} \times \{.05, .4\}$ . To have a realistic range parameter value, we set it to  $0.7R_{max}$ , where  $R_{max} = 2(\frac{2}{\pi^2 \lambda})^{1/2}$  is the maximal range under hard packing limit. In setting the range parameters the unknown intensity  $\lambda$  was crudely replaced by  $\exp(\theta_1)$ , leading to  $R = \{.06, .02\}$  for small and large  $\theta_1$ , respectively. The observation window was set to  $[0, 1]^2$  with additional  $R$ -dilation for border correction.

We set three types of priors for  $\theta$ :

- a) "flat",  $\mu_a = 0^T, S_a = \text{diag}(10^9)$
- b) "tight" around correct values,  $\mu_b = \theta^T, S_b = \text{diag}(1, .01)$  when  $\theta_2 = .05$  and  $S_b = \text{diag}(1, .001)$  when  $\theta_2 = .4$
- c) "tight" around wrong value  $\mu_c = \log(2) + \theta^T$  and  $S_c = S_b$ .

The "tight" wrong priors (c) correspond to stating "expert knowledge" that the true interaction parameter  $\gamma = \exp(\theta_2)$  is between .07 and .19 when in fact  $\gamma = .05$ , and between .65 and 1 when  $\gamma = .4$ . The estimation used the true values of  $R$  parameters as is the common practice in similar experiments.

The same dummy point configurations were used for each fitter, and were provided by the current state-of-the-art Gibbs model inference function `ppm` from the `spatstat`-package (Baddeley and Turner, 2005).

Figure 1 shows the estimated parameter values from 100 simulations with the described Bayesian method (VB) plotted against the standard GLM-fitter `glm` of the R-software as used by the `ppm`-function. Left column, with the "flat" prior, shows that the VB technique with weak priors induces virtually no extra bias to the estimates, confirming the past experiences of good performance of the tangential approximation (Jaakkola and Jordan, 2000).

The middle column, with tight prior around the true value (b), shows that the for low  $< 100$  points intensity the estimates are concentrated strongly around the prior (0-line on the figure). Higher intensity diminishes the prior effect. In the right column we see that the wrong prior (c) effect is likewise strong for small intensity and diminishes when more points are available (cf. + vs ×). The reduction is higher for less interaction (larger  $\theta_2$ ). It is clear that to fully override strong priors an intensity around 1000 is not enough, but for large point patterns (of order  $> 10^4$ ) a dubious expert knowledge would most likely be overridden by the data.




Figure 1: Comparison of the VB logistic regression technique and the standard `glm`-fitter in R. 100 simulations of Strauss model with 4 sets of parameters (see legend). Columns correspond to different priors (see text). Upper and lower row correspond to  $\theta_1$  and  $\theta_2$ , respectively, units are in relative error.

### 3.2 Mucuous cells with trends

This example revisits the Example 5.4 of [Baddeley et al. \(2014\)](#). The data consist of mucous membrane cells of two types, see Figure 2a). The interesting question is whether the two cell types have a different trend component, indicating that the intensities are not proportional. The trend model for each type is a fourth degree polynomial in the vertical coordinate. A Strauss component is set for cross-type interaction, with range fixed to  $R = 0.008$ .

Figure 2b) depicts the posterior pointwise 95%-envelopes for the trends without intercept term, derived using 1000 simulations from the posterior of the 10-dimensional trend parameter vector. The trends are clearly separate, especially at large  $y$ -coordinate locations. A confirmation is given by the Bayes factor 26 against a model with a shared trend.

Figure 2b) answers the same question as Figure 2b) presented by [Baddeley et al. \(2014\)](#), namely what are the pointwise confidence regions for the trends given the data. To derive the regions [Baddeley et al. \(2014\)](#) deploy a parametric bootstrap scheme where the fitted model is repeatedly simulated. Proper simulation of a point process model is in general not trivial, and can in fact be very costly. In our case we need only simulate a 10-dimension Gaussian random variable, a comparatively trivial execution with very little cost.

### 3.3 Inference on an unknown interaction function

For the so-called pairwise interaction Gibbs model family an often used alternative form of the density is

$$\tilde{f}(\omega; \theta) \propto \prod_i \phi_0(\omega_i) \prod_{i < j} \phi_\theta(d_{ij}), \quad d_{ij} := \|\omega_i - \omega_j\|$$

where  $\phi_0 : \mathbb{R}^d \rightarrow \mathbb{R}^+$  is related to the trend ("first order interaction") and  $\phi_\theta : \mathbb{R}^+ \rightarrow \mathbb{R}^+$  is called the (pairwise) interaction function. The connection to exponential form is  $\prod \phi_\theta(d_{ij}) = \exp[\sum_{i < j} \log \phi_\theta(d_{ij})]$ . For example, the

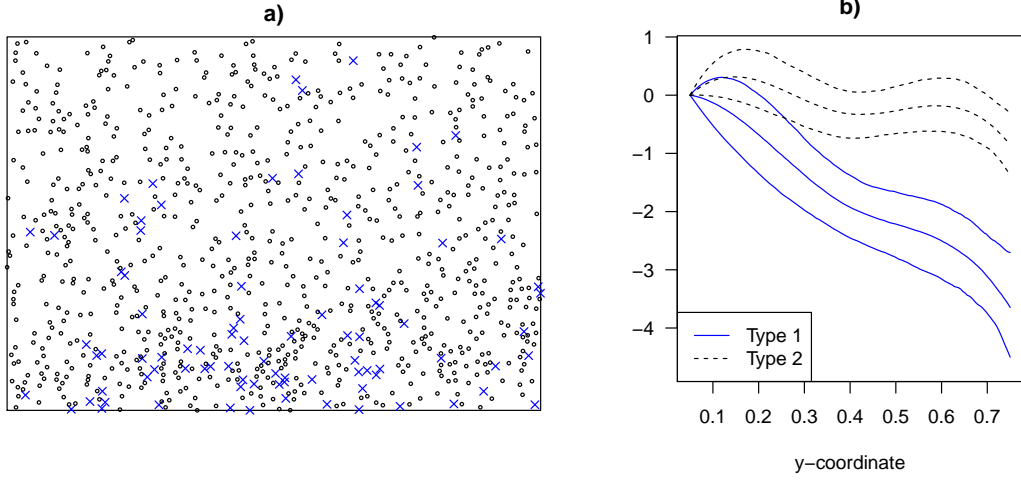


Figure 2: a) Mucosa dataset, window  $[0, 1] \times [0.05, 0.75]$  type 1 =  $\circ$ , type 2 =  $\times$ . b) Posterior pointwise 95%-envelopes for the trends without intercept, as a function of  $y$ -coordinate.

Strauss model belongs to this family with the interaction function  $\phi_\theta(r) = \exp(\theta_2)^{1(r < R)}$ .

If we do not have a good idea of the shape of the interaction function, i.e. a priori we do not set a specific interaction model, we can approximate it by a step function: Given a set of ranges  $0 = r_0 < r_1 < r_2 < \dots < r_K = r_{\max}$ , a function

$$\tilde{\psi}_\theta(r) := \sum_{k=1}^K 1\{r_{k-1} \leq r < r_k\} w_k = \sum_{k=1}^K h_k(r) w_k \quad (6)$$

with some weights  $w_k \in \mathbb{R}$  can be constructed so that it converges to  $\log \phi(r)$  as  $K \rightarrow \infty$ .

With a Gaussian prior for the weights  $\theta_{2:(K+1)} = \mathbf{w}$ , the conditional intensity of the approximating semi-parametric model becomes

$$\lambda(u_i; \omega) = \exp(\text{1st order term} + \mathbf{w}^T \psi_i)$$

where  $\psi_i = [\sum_j h_1(d_{ij}) \dots \sum_j h_K(d_{ij})]^T$  with the coding that we exclude self-interaction and interaction to dummy points, i.e.  $h_k(d_{ij}) = 0$  whenever  $i = j$  or  $j > n$ . The above conditional intensity is that of an exponential model so we are back at the logistic regression situation.

Within the Bayesian framework, we can impose soft constraints on the weights via priors. For example, a typical repulsive interaction function is at  $r = 0$  equal to 0 and for short ranges close to 0, and at  $r_{\max}$  interaction converges to 1 (no interaction) as stability conditions require the interaction to be of finite range. We can set the weights to be conditional on these bordering values, and connect the intermediate weights by imposing off-diagonal covariances in Gaussian process fashion, using some suitable kernel function. Note that we are not modeling the interaction function with a multi-scale family but merely approximating it like [Heikkinen and Penttinen \(1999\)](#) did using MCMC. In their implementation the weights were strictly increasing, but unfortunately it is not possible to formulate such hard conditioning in the prior construction.

In Figure 3a) the blue thick line depicts the interaction function of the [Lennard-Jones \(1924\)](#) model, a model motivated by atomic and molecular interactions. A simulation from the model was generated (not shown) in a window  $[0, 2]^2$  with first order intensity related term  $\theta_1 = \log(200)$ . Overlaid on the true function are LOESS-smoothed samples from the posterior of the step-function (6), along with the posterior mean of  $w_k$ 's. The estimate of the interaction function is quite acceptable, and an estimate of the characteristic repulsion range with certainty interval can be inferred from the posteriors. Via the formula for the Lennard-Jones interaction function's maxima, the mean estimate of the range is .058 with 95% central confidence region  $[\text{.057}, \text{.062}]$ , while the true characteristic range used in the simulation was .06.

As a further thought: The step-function approximation can also be extended to the general basis extension approximation. We already did a similar non-parametric approximation for the unknown trend functions in the mucous example, using polynomial basis functions. Figure 3b) depicts the estimate for the same Lennard-Jones

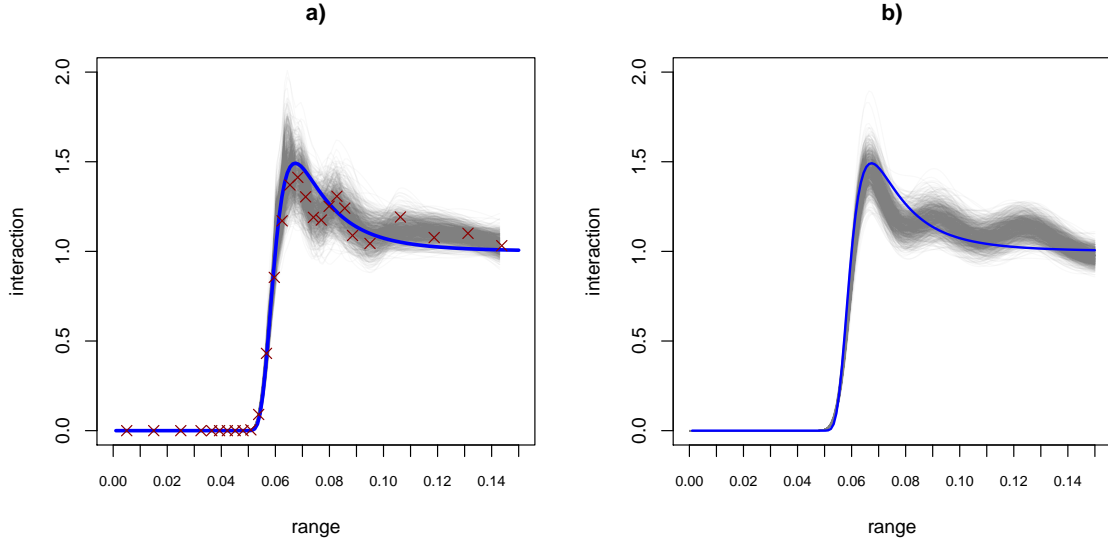


Figure 3: a) Interaction function estimation using a step-function with a smoothing prior. True interaction used in simulation (blue, thick line), 1000 posterior simulations of the step function, LOESS-smoothed (light gray lines), and the posterior means of the step function weights (red  $\times$ ). b) Interaction function estimated using 50 square exponential basis functions.

interaction function using 50 squared-exponential basis functions as  $h_k$ , using the same  $r$ -grid for centers as for the step functions. The results are again quite good, considering no particular information of the original potential's shape is used. Characteristic range estimate is also very good,  $.059 \pm .010$ .

The resolution of the  $r$ -grid and the basis function width, both corresponding to the level of smoothness of the unknown function is an issue that has been discussed heavily in literature and therefore beyond this text.

## 4 Discussion

In this exposition we combined two approximation techniques to get a computationally cheap Bayesian method for analysing exponential family Gibbs point processes. The first of the two approximations, and in all accounts the more important one in point pattern context, draws the connection between the conditional intensity of a process and the logistic regression, as described by [Baddeley et al. \(2014\)](#) and apparently discovered by others during the past decade (personal communications). As it is superior to the most commonly used Poisson approximation to the pseudo-likelihood, the connection is currently the most efficient method for estimating the parameters of these models and is opening interesting doors for point pattern analysis and modeling.

We developed a Bayesian version of the logistic regression method by recalling the second approximation from machine learning literature, a technique that has seen surprisingly little use in the hands of statisticians. The experience on approximating the logistic likelihood using a tangential functions has a good track record in machine learning community, where feasibility, performance and scalability are valued, and the examples shown here only reinforce that optimism.

The shown examples are not particularly new. Bayesian inference for unknown interaction functions and posterior analysis of trend components have been described earlier by e.g. [Heikkinen and Penttinen \(1999\)](#), [Berthelsen and Møller \(2003\)](#) and [Bognar \(2005\)](#). However, most of Bayesian analysis of point pattern data are based on MCMC, which is often time consuming for exponential point process models, involving repeated simulations of the process. The described method avoids all such simulations, thus providing real savings on time and computations. And with a full posterior distribution for the parameters at hand some further computational strain can be alleviated in the data analysis over the non-Bayesian counterparts. In the analysis of the mucuous data, for example, the bootstrap simulations of a bivariate Gibbs point process was replaced by considerably easier simulations of a Gaussian vector.

Hopefully this discussion increases the interest of statistical community on the developments happening in other data-oriented fields with similar computational problems. The emphasis on analysis details might be different, but



the numerical problems are parallel and we should use that connection to everyones advantage.

## 5 Acknowledgements

The author has been financially supported by the Knut and Alice Wallenberg Foundation.

## References

- A. Baddeley and R. Turner. Spatstat: an R package for analyzing spatial point patterns. *Journal of Statistical Software*, 12:1–42, 2005. URL <http://www.jstatsoft.org>.
- A. Baddeley, J. Coeurjolly, E. Rubak, and R. Waagepetersen. Logistic regression for spatial Gibbs point processes. *Biometrika*, 101:377–392, 2014.
- K. Berthelsen and J. Møller. Likelihood and Non-parametric Bayesian MCMC Inference for Spatial Point Processes Based on Perfect Simulation and Path Sampling. *Scandinavian Journal of Statistics*, 30:549–564, 2003.
- M. Bognar. Bayesian inference for spatially inhomogeneous pairwise interacting point processes. *Computational Statistics & Data Analysis*, 49:1–18, 2005. ISSN 01679473.
- A. Gelfand, P. Diggle, M. Fuentes, and P. Guttorp, editors. *Handbook of Spatial Statistics*. CRC Press, 2010.
- J. Heikkinen and A. Penttinen. Bayesian smoothing in the estimation of the pair potential function of Gibbs point processes. *Bernoulli*, 5:1119–1136, 1999. URL <http://projecteuclid.org/euclid.bj/1143122305>.
- T. Jaakkola and M. Jordan. Bayesian parameter estimation via variational methods. *Statistics and Computing*, 10:25–37, 2000. URL <http://www.springerlink.com/index/V651T420415P2312.pdf>.
- J. E. Lennard-Jones. On the Determination of Molecular Fields. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 106:463–477, 1924. ISSN 1364-5021.
- J. Møller and R. Waagepetersen. Modern Statistics for Spatial Point Processes. *Scandinavian Journal of Statistics*, 34:643–684, 2007. ISSN 0303-6898.
- J. Ormerod and M. Wand. Explaining variational approximations. *The American Statistician*, 64:140–153, 2010. URL <http://amstat.tandfonline.com/doi/abs/10.1198/tast.2010.09058>.